UNIVERSITY OF SASKATCHEWAN

Department of Computational Science

CMPT 422.3/835.3

Final Examination

CLOSED BOOK

Time: 3 Hours

April 15, 1994

INSTRUCTIONS:

1. There are five questions on this examination. They may be answered in any order, but all five are to be answered.
2. Read each question carefully and take time to plan your answer. A portion of the marks for each question will be awarded for the organization, clarity, and precision of the answer.
3. Apportion your time according to the indicated mark values.

Marks

10   1.   (a)   Theoretically optimal strategies can be formulated for the allocation of various system resources. Describe in detail, and illustrate by means of a suitable example, the optimal allocation strategy for each of the following:

       (i)   page replacement in a page-on-demand system,
       (ii)   job management under intermittent arrivals.

       In each case describe in full the information required to achieve optimal allocation.

10   (b)   Central to any resource management strategy is an underlying assumption about the characteristics of demands for this resource. Describe the "ideal model" in each of the following and defend or criticize each assumption in a real system setting.

       (i)   Page replacement algorithms RAND, LFU, and LRU.
       (ii)   Scheduling algorithms FCFS, SPT, and FB

5 / 2. (a)   When do page faults occur? Describe the actions taken by an operating system when a page fault occurs.

9   (b)   Consider the following sequence of memory references from a 460-byte program:

10, 11, 104, 170, 73, 309, 185, 245, 246, 434, 458, 364

   (i)   Give the reference string, assuming a page size of 100 bytes.
   (ii)   Compute the total number of page faults for this reference string assuming a memory allocation of 200 bytes and the LRU replacement algorithm.
   (iii)   Compute the total number of page faults for this reference string assuming a memory allocation of 200 bytes and Belady's optimal replacement algorithm (MIN).

6   (c)   Some page replacement strategies result in a resident set whose size varies over time while others result in a constant resident set size. Illustrate by giving an example of a strategy of each type. Which approach makes more sense and why?

10   3. (a)   Monitors and semaphores each provide the means by which process interactions can be controlled in an operating system. Outline the approach taken in each case and discuss the strengths and weaknesses of each approach.

10   (b)   There are two basic approaches to interprocess communication in an operating system, the shared memory approach and the message passing approach. Discuss briefly how each works and give the advantages and disadvantages of each approach.

12   4. (a)   Suppose the head of a moving-head disk with 200 tracks, numbered 0 to 199, is currently serving a request at track 143 and has just finished serving a request at track 125. The following queue of requests (in order of arrival) awaits service:

86, 147, 91, 177, 94, 150, 102, 175, 130

Compute the total head movement needed to satisfy these requests for the following disk scheduling algorithms: FCFS, SSTF, SCAN, LOOK.

8   (b)   Delayed writeback can be used to advantage in both paging and file caching. What is it, why is it advantageous, and what problems might it create in each application?

5. It has been said that an operating system is a community of resources, and that a comprehensive approach to resource management is the only reasonable approach.

10 (a) Identify and discuss a situation in which failing to consider simultaneous resource demands (in either a single process or across a set of processes) can lead directly to performance problems.

10 (b) Identify and discuss a situation in which explicit consideration of simultaneous resource demands (in either a single process or across a set of processes) can lead to improved performance.

THE END

100 marks